

Docket No. AUS920030443US1

**APPARATUS, SYSTEM AND METHOD OF EXECUTING MONOLITHIC
APPLICATION PROGRAMS ON GRID COMPUTING SYSTEMS**

5

BACKGROUND OF THE INVENTION

1. Technical Field:

The present invention is directed to execution of programs on grid computing systems. More specifically, the present invention is directed to a system and method of executing monolithic programs on grid computing systems.

2. Description of Related Art:

With the advent of the Internet and the availability of powerful computers and high speed networks, geographically dispersed computers are beginning to be used as one single unified computing resource. This computing resource is popularly referred to as grid computing (also known as meta-computing, scalable computing, global computing, Internet computing and peer-to-peer computing). In grid computing, resources such as supercomputers, storage systems, data sources etc. which are owned by different organizations are pooled together to solve large-scale computational and data intensive problems in science, engineering and commerce.

To allow for an effective use of grid computing, a grid broker is used. The grid broker performs resource discovery and scheduling of tasks on the distributed grid resources. Thus, if the grid broker receives a request from a computer in New York to process an application program, the grid broker may, based on a resource discovery result, schedule the application program to be processed by a computer in Paris, France.

Docket No. AUS920030443US1

This concept works fine with application programs that are grid-computing compatible or have been designed to be processed by any computer system on a grid computing system. However, many older application programs (i.e., legacy programs) are monolithic. A monolithic application program is a program that does not rely on outside resources and cannot access or provide services to other applications in a dynamic and cooperative manner. An example of such an application program is a program that may link to specific (local) libraries or read from and/or write to specific remote filesystems etc.

Due to the nature of monolithic application programs, they may not run effectively on a grid computing system without modifications. To modify a monolithic program, however, requires that software personnel acquire intimate knowledge of the program. In some instances, this may be quite a time-consuming and work-intensive endeavor, especially, when there is only an object code of the program available and all software personnel who may have designed and developed the program cannot be located.

Consequently, what is needed is an apparatus, system and method of executing a monolithic application program successfully on a grid computing system without modifications.

25

Docket No. AUS920030443US1

SUMMARY OF THE INVENTION

The present invention provides an apparatus, system and method of executing a monolithic application program successfully on a grid computing system. Before the program is executed on the grid computing system, the program is executed on a computer on which the program has previously been successfully executed. During its execution, the program is monitored to collect its runtime information. The runtime information is provided to the grid computing system. With this information the grid computing system is able to successfully execute the program.

Docket No. AUS920030443US1

BRIEF DESCRIPTION OF THE DRAWINGS

The novel features believed characteristic of the invention are set forth in the appended claims. The invention itself, however, as well as a preferred mode of use, further objectives and advantages thereof, will best be understood by reference to the following detailed description of an illustrative embodiment when read in conjunction with the accompanying drawings, wherein:

10 Fig. 1 is an exemplary block diagram illustrating a distributed data processing system according to the present invention.

 Fig. 2 is an exemplary block diagram of a server apparatus according to the present invention.

15 Fig. 3 is an exemplary block diagram of a client apparatus according to the present invention.

 Fig. 4 illustrates a representative grid computing system.

20 Fig. 5 is a flowchart of a process that may be used by the invention.

 Fig. 6 is a flowchart of a first process that may be used by the invention.

 Fig. 7 is a flowchart of a second process that may be used by the invention.

25

Docket No. AUS920030443US1

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

With reference now to the figures wherein like reference numerals refer to like parts, Fig. 1 depicts a pictorial representation of a network of data processing systems in which the present invention may be implemented. Network data processing system 100 is a network of computers in which the present invention may be implemented. Network data processing system 100 contains a network 102, which is the medium used to provide communications links between various devices and computers connected together within network data processing system 100. Network 102 may include connections, such as wire, wireless communication links, or fiber optic cables.

In the depicted example, server 104 is connected to network 102 along with storage unit 106. In addition, clients 108, 110, and 112 are connected to network 102. These clients 108, 110, and 112 may be, for example, personal computers or network computers. In the depicted example, server 104 provides data, such as boot files, operating system images, and applications to clients 108, 110 and 112. Clients 108, 110 and 112 are clients to server 104. Network data processing system 100 may include additional servers, clients, and other devices not shown. In the depicted example, network data processing system 100 is the Internet with network 102 representing a worldwide collection of networks and gateways that use the TCP/IP suite of protocols to communicate with one another. At the heart of the Internet is a backbone of high-speed data communication lines between major nodes or host computers, consisting of thousands of commercial, government, educational and other computer systems that route data and

Docket No. AUS920030443US1

messages. Of course, network data processing system 100 also may be implemented as a number of different types of networks, such as for example, an intranet, a local area network (LAN), or a wide area network (WAN). Fig. 1 is
5 intended as an example, and not as an architectural limitation for the present invention.

Referring to Fig. 2, a block diagram of a data processing system that may be implemented as a server, such as server 104 in Fig. 1, is depicted in accordance with a
10 preferred embodiment of the present invention. Data processing system 200 may be a symmetric multiprocessor (SMP) system including a plurality of processors 202 and 204 connected to system bus 206. Alternatively, a single processor system may be employed. Also connected to system
15 bus 206 is memory controller/cache 208, which provides an interface to local memory 209. I/O bus bridge 210 is connected to system bus 206 and provides an interface to I/O bus 212. Memory controller/cache 208 and I/O bus bridge 210 may be integrated as depicted.

20 Peripheral component interconnect (PCI) bus bridge 214 connected to I/O bus 212 provides an interface to PCI local bus 216. A number of modems may be connected to PCI local bus 216. Typical PCI bus implementations will support four PCI expansion slots or add-in connectors. Communications
25 links to network computers 108, 110 and 112 in Fig. 1 may be provided through modem 218 and network adapter 220 connected to PCI local bus 216 through add-in boards. Additional PCI bus bridges 222 and 224 provide interfaces for additional PCI local buses 226 and 228, from which additional modems or
30 network adapters may be supported. In this manner, data processing system 200 allows connections to multiple network computers. A memory-mapped graphics adapter 230 and hard

Docket No. AUS920030443US1

disk 232 may also be connected to I/O bus 212 as depicted, either directly or indirectly.

Those of ordinary skill in the art will appreciate that the hardware depicted in Fig. 2 may vary. For example, 5 other peripheral devices, such as optical disk drives and the like, also may be used in addition to or in place of the hardware depicted. The depicted example is not meant to imply architectural limitations with respect to the present invention.

10 The data processing system depicted in Fig. 2 may be, for example, an IBM e-Server pSeries system, a product of International Business Machines Corporation in Armonk, New York, running the Advanced Interactive Executive (AIX) operating system or LINUX operating system.

15 With reference now to Fig. 3, a block diagram illustrating a data processing system is depicted in which the present invention may be implemented. Data processing system 300 is an example of a client computer. Data processing system 300 employs a peripheral component
20 interconnect (PCI) local bus architecture. Although the depicted example employs a PCI bus, other bus architectures such as Accelerated Graphics Port (AGP) and Industry Standard Architecture (ISA) may be used. Processor 302 and main memory 304 are connected to PCI local bus 306 through
25 PCI bridge 308. PCI bridge 308 also may include an integrated memory controller and cache memory for processor 302. Additional connections to PCI local bus 306 may be made through direct component interconnection or through add-in boards. In the depicted example, local area network
30 (LAN) adapter 310, SCSI host bus adapter 312, and expansion bus interface 314 are connected to PCI local bus 306 by direct component connection. In contrast, audio adapter

Docket No. AUS920030443US1

316, graphics adapter 318, and audio/video adapter 319 are connected to PCI local bus 306 by add-in boards inserted into expansion slots. Expansion bus interface 314 provides a connection for a keyboard and mouse adapter 320, modem 5 322, and additional memory 324. Small computer system interface (SCSI) host bus adapter 312 provides a connection for hard disk drive 326, tape drive 328, and CD-ROM drive 330. Typical PCI local bus implementations will support three or four PCI expansion slots or add-in connectors.

10 An operating system runs on processor 302 and is used to coordinate and provide control of various components within data processing system 300 in Fig. 3. The operating system may be a commercially available operating system, such as Windows® XP, which is available from Microsoft Corporation. An object oriented programming system such as 15 Java may run in conjunction with the operating system and provide calls to the operating system from Java programs or applications executing on data processing system 300. "Java" is a trademark of Sun Microsystems, Inc. 20 Instructions for the operating system, the object-oriented operating system, and applications or programs are located on storage devices, such as hard disk drive 326, and may be loaded into main memory 304 for execution by processor 302.

Those of ordinary skill in the art will appreciate that 25 the hardware in Fig. 3 may vary depending on the implementation. Other internal hardware or peripheral devices, such as flash ROM (or equivalent nonvolatile memory) or optical disk drives and the like, may be used in addition to or in place of the hardware depicted in Fig. 3. 30 Also, the processes of the present invention may be applied to a multiprocessor data processing system.

Docket No. AUS920030443US1

As another example, data processing system 300 may be a stand-alone system configured to be bootable without relying on some type of network communication interface, whether or not data processing system 300 comprises some type of network communication interface. As a further example, data processing system 300 may be a Personal Digital Assistant (PDA) device, which is configured with ROM and/or flash ROM in order to provide non-volatile memory for storing operating system files and/or user-generated data.

The depicted example in Fig. 3 and above-described examples are not meant to imply architectural limitations. For example, data processing system 300 may also be a notebook computer or hand held computer in addition to taking the form of a PDA. Data processing system 300 also may be a kiosk or a Web appliance.

The present invention provides an apparatus, system and method of having a monolithic application program successfully execute on a grid computing system without modifications. The invention may be local to client systems 108, 110 and 112 of Fig. 1 or to the server 104 or to both the server 104 and clients 108, 110 and 112. Further, the present invention may reside on any data storage medium (i.e., floppy disk, compact disk, hard disk, ROM, RAM, etc.) used by a computer system.

Fig. 4 illustrates a representative grid computing system. The grid computing system contains a client user 402 with an agent 404 running thereon. The grid computing system further contains a server 410 connected to a data repository 412, a grid trade center 420, a plurality of resources providers 450 with agent 452 running thereon, a job management 460, accounting system 470 and billing system 472.

Docket No. AUS920030443US1

The server 410, repository 412 and grid trade center 420 for all intent and purpose form the grid broker 430. Particularly, when user 402 needs to have a program processed by the grid computing system, it will so notify
5 server 410 using agent 404. Upon receiving the notification, the server 410 will consult data repository 412 to determine which resources providers are available to process the program. In this particular example, the server 410 indicates that resources providers 450 are able to
10 process the program. Using this information, the client user 402 then establishes accounts with resources providers 450 through grid trade center 420. The client user 402 then notifies the server 410 which will schedule the program for processing.

15 Note that the above scenario is one of a plurality of scenarios that may used by a grid computing system to schedule tasks. In other scenarios, the broker may lease grid resources and services depending on their capability, cost and availability driven by user quality of service
20 (QoS) requirements. In those instances, a user may set a deadline by which results of a task are needed. The broker then has to find the cheapest computational resources available on the grid and use them such that the user deadline is met and the cost of computation is kept to a
25 minimum. Thus, the scenario outlined above is used for illustrative purposes only.

In any case, the client user may manage the processing of the program through job management system 460. The job management system 460 may be an I/O server where a task
30 script (not shown) containing executing commands may be stored. Note that the server 410 may deploy agents 452 on the resources providers 450 in order to manage the

Docket No. AUS920030443US1

scheduling and execution of assigned tasks. After execution of a task, the leasing rate and execution time may be sent to accounting system 470 where a bill may be sent to client user 402 through billing system 472.

5 As mentioned in the Description of Related Art Section, monolithic application programs may not be able to be processed by a grid computing system. For example, suppose an application program uses a library file "A" that is local to a computer system on which it normally executes.
10 Suppose, further, that the program reads from and writes to a file system "B". Then, in order for the program to successfully execute, file "A" must be loaded in the system memory of the computer system executing the program or linked to (depending on the implementation) the executing
15 program. In addition, file system "B" must be mounted on the computer system in order for the program to be able to read from and write to it. Since when the program is being executed on a grid computing system it is most likely being executed by a computer system that does not have file "A"
20 stored thereon and/or file system "B" mounted thereon, a runtime error will most probably ensue.

 The present invention provides a grid application survey (GAS) which may be used to enable a monolithic application program to be processed on a grid computing
25 system without runtime errors. According to the invention, before a monolithic application program is ever processed on a grid, it is executed on a computer system on which it has previously successfully executed in conjunction with the GAS. The GAS monitors the execution of the program and
30 records its runtime requirements. Specifically, the GAS monitors the computer's logical file system for accesses to remote file systems such as networked file systems (NFSs)

Docket No. AUS920030443US1

etc. The GAS then records the necessary mount information from the logical file system's mount table.

A logical file system provides functions required to map system call requests to virtual file system requests.
5 The logical file system is responsible for resolution of file names and file descriptors. It tracks all open files in the system using a file table.

An NFS is used to allow hosts to share files across a network. The NFS operates on a client-server basis using
10 network messages to transfer requests and data between clients and servers. Provided that appropriate drivers are installed, several different sorts of file systems may be mounted. These file systems may include normal Unix file systems, NFSs, "High Sierra" file systems (i.e., on a CDROM)
15 and MSDOS file system (known as PCFS) on a floppy disc. Such systems are integrated into the Unix virtual file system using mount commands.

The GAS also monitors and records files that are linked to the program and/or loaded into memory at runtime. For
20 example, if file "A" of the example above, was loaded in memory of the computer system executing the program or linked to the program, the GAS will note that information down.

The GAS uses the information gathered to generate a
25 setup script that contains all necessary mount information as well as information regarding all files that are linked to and/or loaded in memory. The script may then be incorporated into the task script in the job management system 460. Thus, when the monolithic program is being
30 executed on the grid, the script may transparently execute proper commands to mount the proper file systems on the

Docket No. AUS920030443US1

resources provider executing the program as well as to link the proper files to the program and/or load them in memory.

Again to continue with the example above, the setup script may have proper commands to load file "A" in memory
5 of the resources provider 450 executing the program or facilitate the executing program to link to file "A". Likewise, the setup script may have commands to mount file system "B" onto the resources provider 450 thereby allowing the program to successfully read from or write to the file
10 system.

Alternatively, the script may be executed before the monolithic program is executed on the grid. In this case, the script may be executed by one of the resources providers 450 that is going to execute the program. During its
15 execution, the script may mount file system "B", and link to and/or load file "A" to allow for the successful execution of the program. Upon completion of its execution, the script may instantiate the program.

Fig. 5 is a flowchart of a process that may be used by
20 the invention. The process starts when the program is being executed on the native computer system to gather up required runtime information. If the program is not being executed on the native computer system to gather up the information, the process ends. Otherwise, the information is collected
25 by monitoring the runtime requirements of the program (steps 500, 505, 510 and 515).

Fig. 6 is a flowchart of a first process that may be used by the invention. The process starts when the program is being executed on the grid computing system. As
30 mentioned before, the GAS (i.e., the process) will be running on the job management system 460. Consequently, the GAS will know on which one of the resources providers 450

Docket No. AUS920030443US1

the program is being executed. With this information, the process will mount all required file systems onto the proper resources provider 450. The process will further ensure that all files that need to be linked to the executing
5 program as well as files that need to be loaded into the working memory of the resources provider 450 are either linked to the program or loaded in memory before the process ends (steps 600, 605, 610, 615 and 620).

Fig. 7 is a flowchart of a second process that may be
10 used by the invention. The process starts when a resources provider 450 is to execute the program. In this case, the process will be pre-pended to the program. Thus, the process will run first. By so doing, the process will mount all required file systems onto the resources provider 450.
15 In addition, the process will link to and/or load all requisite files, I/O files as well as library files, in memory. Before exiting, the process will instantiate the program (steps 700, 705, 710, 715 and 720).

The description of the present invention has been
20 presented for purposes of illustration and description, and is not intended to be exhaustive or limited to the invention in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art. The embodiment was chosen and described in order to best explain
25 the principles of the invention, the practical application, and to enable others of ordinary skill in the art to understand the invention for various embodiments with various modifications as are suited to the particular use contemplated.